# *BABAR*'s Experience with the Preservation of Data and Analysis Capabilities

*Marcus* Ebert[1,*], *Michael* Roney[1], and *Randall* Sobie[2]

[1]Department of Physics and Astronomy, University of Victoria, BC, Canada
[2]Institute of Particle Physics, University of Victoria, BC, Canada

**Abstract.** The *BABAR* experiment collected electron-positron collisions at the SLAC National Accelerator Laboratory (SLAC) from 1999-2008. Although data taking has stopped 15 years ago, the collaboration is still actively doing data analyses, publishing results, and giving presentations at international conferences. Special considerations were needed to do analyses using a computing environment that was developed decades ago. A framework is required that preserves the data, data access, and the capability of doing analyses using a well defined and preserved environment. Also, *BABAR*'s support by SLAC ended at the beginning of 2021. Fortunately, the High Energy Physics Research Computing group at the University of Victoria (UVic), Canada, offered to provide the new home for the main *BABAR* computing infrastructure, the Grid Computing Centre Karlsruhe offered to host all data for access by analyses running at UVic, and CERN and the IN2P3 Computing Centre offered to store a backup of all data. This paper presents what was done at *BABAR* to preserve the data and analysis capabilities and what was needed to move the whole computing infrastructure, including collaboration tools and data files, away from SLAC. It will be shown how *BABAR* preserved the ability to continue to do data analyses and also have a working collaboration tools infrastructure. This paper will describe on *BABAR*'s experience with such a big change in its infrastructure and what was learned from it, which may be useful to other experiments which are interested in long term analysis support and data preservation in general.

## 1 Introduction

The legacy *BABAR* computing system was located at SLAC National Accelerator Laboratory (SLAC)[1] where the experiment took place from 1999 to 2008 when data collection ended[2]. Over time, *BABAR*'s computing system grew into the general IT infrastructure of SLAC and became integrated in central databases and systems. The support for *BABAR* at SLAC was extended to the beginning of 2021. At that time, tools and systems used in *BABAR*'s computing framework relied fully on the SLAC systems. This applies to the analysis system where the framework relies on local path names and the local batch system, as well as to collaboration tools, like web pages, calendars, the analysis and member database, and forums for information exchange. In addition, all data files were stored on tape and accessed using a

*e-mail: mebert@uvic.ca

[1]https://www6.slac.stanford.edu

[2]https://www-public.slac.stanford.edu/babar/

local xrootd [1] cluster at SLAC. To be able to continue analyzing the data and have a working collaboration, all of these different pieces of the computing system needed to be migrated to another location. This includes the collision data collected with the *BABAR* detector as well as Monte Carlo simulation of collision events, the full analysis software framework, the documentation of the full framework and data file content, as well as all of the collaboration tools.

## 2 Preservation of the physics data and its metadata

The key part of a collider experiment is the data collected with the detector as well as Monte Carlo simulations of collision events of interest. Both will be referred to as "data files" in this paper since both have the same structure and are needed to do a data analysis. *BABAR* collected raw data with a combined size of about 700TB. However this data cannot be used directly in an analysis but needs to be processed with one of the software releases and saved in a format usable for data analyses. The derived data files, processed with the latest software release, have a size of a few GBs per file and a combined size of about 1.3PB, and are stored using the root file format [2]. To access data files, xrootd is used and the *BABAR* analysis framework has xrootd support built in to access files remotely through streaming.

### 2.1 Data preservation for active access

To preserve the data at a new location, all files needed to be copied and the metadata system had to be updated to reflect the new location of files. *BABAR*'s metadata is stored in a MariaDB database [3] which contains information about all data files, like the file checksum, the date when the data in a file was collected/created, the detector conditions needed to analyze events stored in a file, to which data sets a file belongs, as well as on which site the file is physically located. *BABAR* used different computing sites in the past and had this concept built into the metadata database. Each computing site used by *BABAR* in the past was usable only for local analyses resulting in the requirement to have a local metadata database as well as a local xrootd system for the data access of local analysis jobs. The Grid Computing Centre Karlsruhe (GridKa)[3] is one of the computing sites used by *BABAR* in the past and still had a local database and xrootd installation for *BABAR*. GridKa's Scientific Data Management department offered to have the data files stored at GridKa, as well as to make the data files accessible via its *BABAR* xrootd installation. GridKa also agreed to keep the metadata database running and accessible.

All files were copied from SLAC to GridKa, using bbcp [4]. To automate the process, scripts were written to query the database for every file locally at SLAC, copy the file from SLAC to GridKa, make sure the checksum of a transferred file is correct, and to update the metadata database at GridKa to mark a file as locally available. The whole process is relatively simple and is limited mainly by the I/O required to move files from tape to disk, to calculate checksums and to transfer all files.

### 2.2 Data preservation for backup purpose

In the past, *BABAR* stored data files at SLAC and also on other sites. At the IN2P3 Computing Centre (CC-IN2P3)[4], data files were stored as a backup in case anything would happen to the data at SLAC. This kind of data storage does not allow for data access by user analyses, but

[3] `https://www.scc.kit.edu/en/research/gridka.php`
[4] `https://cc.in2p3.fr/en/`

it can be used to restore lost data at GridKa. Although an old memorandum of understanding (MoU) with CC-IN2P3 expired, all data files were still stored and a new MoU was signed extending the agreement to at least the end of 2025. Data files created in the last years, which were not available at CC-IN2P3 yet, needed to be copied over from GridKa.

In addition, through the "Data Preservation in High Energy Physics" collaboration (DPHEP)[5], CERN offered to also host a copy of the *BABAR* data for backup purposes. A xrootd space was setup at CERN and all files were copied via xrdcp from GridKa to CERN.

## 3 Analysis framework preservation

To preserve the ability to analyze the physics data, not only the software itself needs to be preserved but also the run time environment. The *BABAR* software is written in C++ and can only compile on a 32bit system. While all of the *BABAR* software code is in its own directory structure, it also depends on tools and versions of those tools that were common more than a decade ago. Other than C++, languages like tcl and php are used within the software stack.

Since it was expected that the manpower within *BABAR* would decline, it was decided to freeze the software framework and the environment in 2012. The environment was fixed to the then-current Scientific Linux 6.3 as Operating System (OS) including its security patch state. To achieve that, *BABAR* switched from using bare-metal worker nodes to a virtualized environment. The OS was installed within a qemu Virtual Machine (VM) and the software directories mounted via NFS. The whole system was then isolated from the rest of the network to account for security issues that may arise due to not applying any security patches anymore. To submit analysis jobs, LSF and TORQUE/MAUI were used as batch system at SLAC. The old system used at SLAC is described in [5].

### 3.1 Preservation of the analysis framework

The preservation of the analysis framework was made possible by using VM images. A cloud computing system based on Openstack using the existing VM image was setup at the University of Victoria (UVic). In addition, local NFS server were setup to provide the software infrastructure and home directories. The High Energy Physics Research Computing (HEPRC)[6] group at UVic uses cloudscheduler to manage VMs on Openstack clouds on demand depending on jobs in an HTCondor queue [6][7]. To make use of the same system, HTCondor was adapted as batch system and wrapper scripts written to translate LSF/TORQUE/MAUI commands to HTCondor commands. The status page of cloudscheduler used for *BABAR* is shown in Fig.1.

While the preservation of the software environment itself was not too complicated due to having all software in a well defined directory structure, issues were found when using the new installation. In addition to the source code for all of the *BABAR* software, all binaries and libraries were precompiled to allow users to run binaries as needed as well as to easily link new code into the framework by using the existing libraries. When trying to run the old, dynamically linked, binaries on the new system, it was found that few of the them were missing libraries. One of the reason was that when including libraries absolute path names were used starting with "/afs/slac.stanford.edu/...". To get around that issue, the old AFS structure had to be replicated in NFS.

Another reason for failure was that the path for some libraries pointed to the home directory of users. Users sometimes used their own login to do production tasks instead of using

[5] https://dphep.web.cern.ch/
[6] http://heprc.phys.uvic.ca/

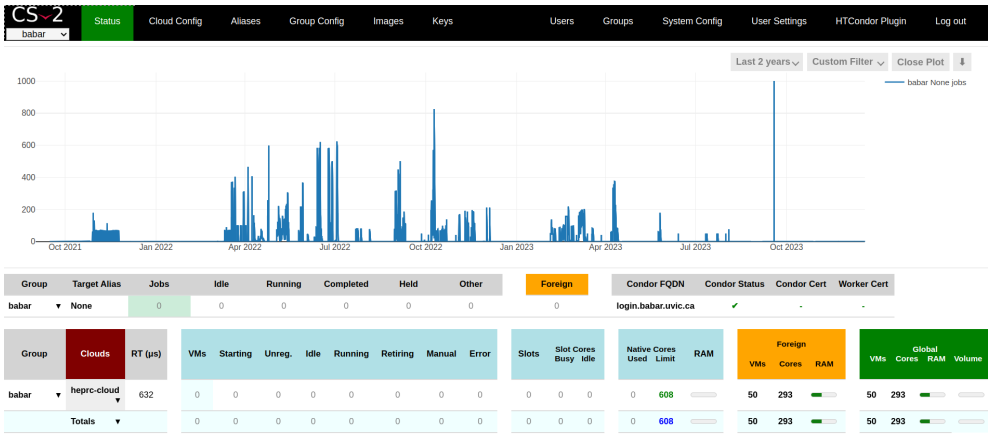


**Figure 1.** Overview of the resource usage by *BABAR* jobs on the new UVic analysis system using cloudscheduler.

specific production accounts and did testing of new code in their own home directories where they build patched binaries. Those binaries got then copied over to the central software area and run fine as long as the user's home directory was available. However, since that was not available outside of SLAC, the binaries failed to execute on the new system. Fortunately, that was found before the old system was shut down and additional directories needed could simply be copied over too.

### 3.2 *BABAR*-To-Go

Since the OS environment needed for the *BABAR* analysis framework was available as a VM image, this image was also made available as a standalone VM together with a second image containing the latest analysis software release. This standalone VM can be downloaded by anyone in *BABAR* and gives the ability to do simple data processing or code development on their own machines. *BABAR*-To-Go also has full access to the data stored at GridKa and gives users a simple way of testing new analysis ideas.

## 4 Documentation preservation

In addition to preserving the physics data files and the analysis framework, the documentation also needs to be preserved. This is especially important when analyses should be possible long term since experts move on to other experiments and expert knowledge within active *BABAR* users will decline over time. A good documentation framework that can easily be updated and is easily available to users will become the only source of how to run an analyses at one point in an experiments life time when aiming for data analysis possibility long after data taking stopped.

### 4.1 Web based documentation

Traditionally, documentations were available as web pages using plain html. All web pages were stored within a well defined directory structure. To make changes to those web pages,

one needed to be logged into a local Linux system and have the right permissions to write to those files. Since everyone in *BABAR* had a SLAC account, that was an easy way to document anything within the collaboration. In addition, page permissions could be set to have specific pages visible to the public while others were only visible to members of the collaboration.

When moving the documentation away from SLAC, the users do not have a local account on any new system and therefore this kind of documentation could no longer remain active. Instead, the content was frozen and made available as static pages for archival purpose and visible to the collaboration only. An additional single public information web page was made available with mostly static content[7]. It is maintained by the local HEP-RC group.

### 4.2 Dynamic web pages

In addition to static web pages which need to be manually updated when needed, some pages also displayed content dynamically via database queries. This was true for most information that needed to be stored in databases, like the list of collaboration members and their details, conference talks given, and so on. Unfortunately, those pages queried an Oracle database within SLAC which is not reachable from outside of the local network. On the other hand, when moving away from SLAC also the databases would no longer be updated. Therefore it was decided to replace those dynamic pages with static pages reflecting the state of the database content at the time the new system was setup. More information about the database replacement can be found in section 5.

### 4.3 Wiki based documentation

A Wiki [8] was setup in 2012 while there was still a larger group of experts around who were able to create the initial documentation in the Wiki system. In addition, everyone within the collaboration has full access to the Wiki and can make changes to it via any web browser. This proved to be a good solution and when moving away from SLAC, the Wiki became the only active documentation. All content of the Wiki is stored in a local MariaDB database which made the move to UVic easy. A new server was setup and the database dumps from the old Wiki used to recreated it locally. As before, all members of the collaboration got full access to its content. This way it can always be kept up to date as long as analyses need to be performed. Permission can also easily be changed in case the collaboration decides one day to make all of its documentation fully public.

As before with the software, unfortunately, links on pages not always used relative names but full URLs. While that was not an issue for the static html pages where the URLs could be changed via simple text processing, for the Wiki content it was more of an issue. Fortunately, SLAC IT agreed to setup a redirect for any access attempt to the old pages at SLAC, redirecting those requests to the new server. In addition, all collaboration members were asked to make manual changes to the Wiki content whenever they find a link pointing to the old location.

## 5 Collaboration tools

This section will cover tools that a collaboration needs to function, like a members list, systems to manage active analyses and bring those to publication, a meeting organizer, mailing lists and a forum to exchange information. Many of the traditionally used systems became

[7] `https://babar.heprc.uvic.ca`

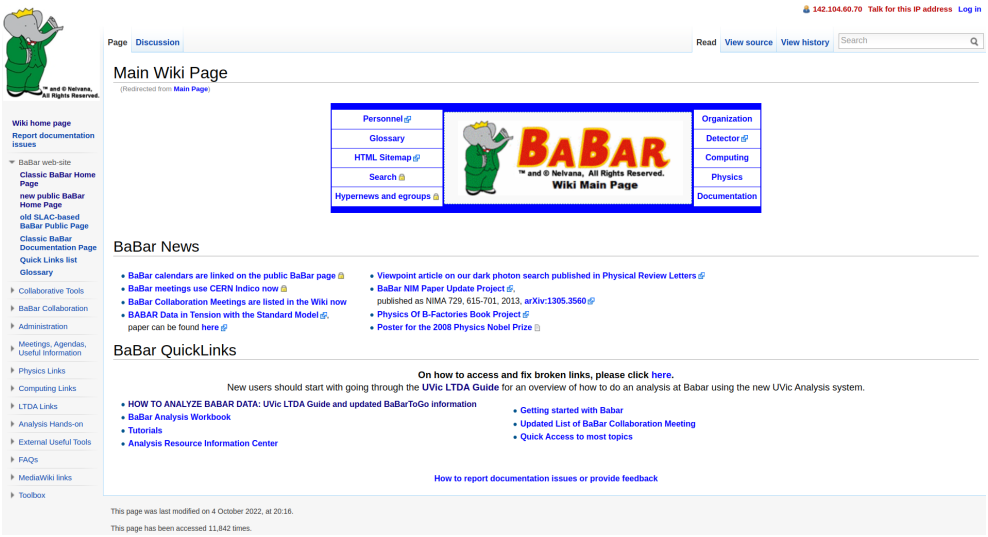


**Figure 2.** Current *BABAR* Wiki main page after moving to UVic.

fully integrated into the SLAC IT infrastructure and cannot easily be ported to another place. Therefore replacements are introduced were needed.

SLAC based mailing lists had to be abandoned and new mailing lists were set up at Caltech; one for *BABAR* associates and one for full *BABAR* members which both also function as member list replacement. In addition, Caltech mailing lists to reach the *BABAR* management team as well as to reach the *BABAR* Speakers Bureau were set up.

CERN not only agreed to store a backup of *BABAR*'s data but also to provide other resources. For organizing meetings, CERN's INDICO system[9] was adapted as well as the CERN egroups for information exchange, discussions as wells as for analysis reviews. The CERN egroups replace the old HyperNews (HN) system that was used before. The old HN system has also been moved to UVic, was made read-only and kept as a historic source of information since people used to discuss all kind of topics and issues with analyses as well as software there.

Other tools were moved to use the Google infrastructure. The Oracle database based analysis documentation storage got replaced by using shared Google Drive folders, utilizing separate folders for each analysis. In addition, other folders were created in Google Drive to have document storage space for the *BABAR* Speakers Bureau as well as the Publication Board. Google sheets and documents are used for different purposes and stored within those folders.

For very long term archival, all analysis documents were moved to INSPIRE[8] who kindly provided a space for *BABAR*. This includes documents submitted for publication as well as *BABAR* internal supporting documents and more detailed information about the analysis and its authors. The backup to INSPIRE is ongoing and documents of newly finished analyses will be moved to INSPIRE.

[8] https://inspirehep.net/

## 6 Conclusion

To preserve the ability to do analyses long after an experiment finished data taking, data files, the whole analysis framework, as well as the documentation need to be preserved and made available. Data preservation and open access to data files alone does not make it possible to analyze data in the future. In addition, if the collaboration still needs to function at this stage, also collaboration tools need to be in place.

*BABAR* went through the process of data and analysis preservation and successfully setup a new infrastructure that will allow performing data analyses in the future. *BABAR*'s data files are preserved, for backup and as well as for active user access, and a new infrastructure for the analysis framework and its documentation are setup at UVic. New collaboration tools were adapted where it made sense or where it was simply not possible to setup the same system outside of SLAC.

Issues that were found in this transition came mostly from the choice of tools in the past, e.g. Oracle database vs freely available open source alternatives. Services that were needed for the analysis framework to function properly were also run using different local Linux accounts instead of having all centrally run using production accounts. This includes the patching and compiling of code resulting in dynamically linked binaries expecting libraries in user's home directories, as well as cronjobs run by users. Other issues that were found came from the fact how code was written, e.g. using absolute SLAC specific paths instead of relative ones where possible. In general, planning already in the early stages of an experiment for data and analysis preservation and choosing tools that can easily be ported to other places will help experiments with future preservation efforts.

## 7 Acknowledgement

Without the help of many different sites, groups and organizations, none of *BABAR*'s effort and success in the preservation of data and analysis capabilities would have been possible. We wish to thank:

- GridKa for providing storage space for all of *BABAR*'s data in a way that it can be accessed by active analyses, as well as for hosting the metadata database needed for data analyses,
- CERN and DPHEP for providing backup space for all of *BABAR*'s data as well as for providing access to important tools like INDICO and egroups,
- CC-IN2P3 for continuing to provide backup space for all of *BABAR*'s data,
- INSPIRE for providing space for analysis document archival, and
- UVic's HEP-RC group for providing and managing the new *BABAR* analysis facility as well as the documentation infrastructure.

## References

[1] *XRootD*, accessed: Sept. 2023, `https://xrootd.slac.stanford.edu`

[2] *The root file format*, accessed: Sept. 2023, `https://root.cern.ch/root/htmldoc/guides/users-guide/InputOutput.html`

[3] *MariaDB*, accessed: Sept. 2023, `https://mariadb.org`

[4] *The bbcp file copy tool*, accessed: Sept. 2023, `https://www.slac.stanford.edu/~abh/bbcp/`

[5] *DPHEP presentation of the BABAR LTDA system at SLAC*, accessed: Sept. 2023, `https://indico.cern.ch/event/209688/contributions/1501412/`

[6] F. Berghaus, K. Casteels, C. Driemel, M. Ebert, F.F. Galindo, C. Leavett-Brown, D. MacDonell, M. Paterson, R. Seuster, R.J. Sobie et al., Computing and Software for Big Science **4**, 4 (2020)

[7] *Cloudscheduler software*, accessed: Sept. 2023, `https://github.com/hep-gc/cloudscheduler`

[8] *The MediaWiki software*, accessed: Sept. 2023, `https://www.mediawiki.org/`

[9] *Indico event and meeting organisation software*, accessed: Sept. 2023, `https://getindico.io/`